Broad Context Language Modeling as Reading Comprehension

Zewei Chu[1] **Hai Wang**[2] **Kevin Gimpel**[2] **David McAllester**[2]

[1]University of Chicago, Chicago, IL, 60637, USA
[2]Toyota Technological Institute at Chicago, Chicago, IL, 60637, USA

zeweichu@uchicago.edu, {haiwang,kgimpel,mcallester}@ttic.edu

arXiv:1610.08431v2 [cs.CL] 19 Dec 2016

Abstract

Progress in text understanding has been driven by large datasets that test particular capabilities, like recent datasets for reading comprehension (Hermann et al., 2015). We focus here on the LAMBADA dataset (Paperno et al., 2016), a word prediction task requiring broader context than the immediate sentence. We view LAMBADA as a reading comprehension problem and apply comprehension models based on neural networks. Though these models are constrained to choose a word from the context, they improve the state of the art on LAMBADA from 7.3% to 49%. We analyze 100 instances, finding that neural network readers perform well in cases that involve selecting a name from the context based on dialogue or discourse cues but struggle when coreference resolution or external knowledge is needed.

1 Introduction

The LAMBADA dataset (Paperno et al., 2016) was designed by identifying word prediction tasks that require broad context. Each instance is drawn from the BookCorpus (Zhu et al., 2015) and consists of a passage of several sentences where the task is to predict the last word of the last sentence. The instances are manually filtered to find cases that are guessable by humans when given the larger context but not when only given the last sentence. The expense of this manual filtering has limited the dataset to only about 10,000 instances which are viewed as development and test data. The training data is taken to be books in the corpus other than those from which the evaluation passages were extracted.

Paperno et al. (2016) provide baseline results with popular language models and neural network architectures; all achieve zero percent accuracy. The best accuracy is 7.3% obtained by randomly choosing a capitalized word from the passage.

Our approach is based on the observation that in 83% of instances the answer appears in the context. We exploit this in two ways. First, we automatically construct a large training set of 1.8 million instances by simply selecting passages where the answer occurs in the context. Second, we treat the problem as a reading comprehension task similar to the CNN/Daily Mail datasets introduced by Hermann et al. (2015), the Children's Book Test (CBT) of Hill et al. (2016), and the Who-did-What dataset of Onishi et al. (2016). We show that standard models for reading comprehension, trained on our automatically generated training set, improve the state of the art on the LAMBADA test set from 7.3% to 49.0%. This is in spite of the fact that these models fail on the 17% of instances in which the answer is not in the context.

We also perform a manual analysis of the LAMBADA task, provide an estimate of human performance, and categorize the instances in terms of the phenomena they test. We find that the comprehension models perform best on instances that require selecting a name from the context based on dialogue or discourse cues, but struggle when required to do coreference resolution or when external knowledge could help in choosing the answer.

2 Methods

We now describe the models that we employ for the LAMBADA task (Section 2.1) as well as our dataset construction procedure (Section 2.2).

2.1 Neural Readers

Hermann et al. (2015) developed the CNN/Daily Mail comprehension tasks and introduced ques-

tion answering models based on neural networks. Many others have been developed since (Kadlec et al., 2016; Chen et al., 2016; Dhingra et al., 2016; Cui et al., 2016). We refer to these models as "neural readers". While a detailed survey is beyond our scope, we briefly summarize the neural readers used in our experiments. We use **d** for the context word sequence, **q** for the question (with a blank to be filled), $\mathcal{A}$ for the candidate answer list, and $\mathcal{V}$ for the vocabulary. We describe neural readers in terms of three components:

1. **Embedding and Encoding**: Each word in **d** and **q** is mapped into a v-dimensional vector via an embedding function $e(w) \in \mathbb{R}^v, w \in \mathbf{d} \cup \mathbf{q}$.[1] The embedded context and question vector sequences are then processed by bidirectional recurrent neural networks (RNN) which compute hidden vectors for each position in each sequence: $h = BiRNN(\boldsymbol{\theta}_d, e(\mathbf{d}))$ and $g = BiRNN(\boldsymbol{\theta}_q, e(\mathbf{q}))$, where $\boldsymbol{\theta}_d$ and $\boldsymbol{\theta}_q$ are BiRNN parameters. They use either gated recurrent units (Cho et al., 2014) or long short-term memory (Hochreiter and Schmidhuber, 1997).

2. **Attention**: Then readers compute attention weights on positions of h and g. In general, we define $\alpha_{ij} = att(h_i, g_j)$, where i and j range over positions in h and g. The att function is typically an inner product or a bilinear function, though others are also used. Most models only use a subset of g, namely the final hidden vector of the forward RNN concatenated with the final vector of the backward RNN. Various other modifications have been proposed (Cui et al., 2016; Dhingra et al., 2016).

3. **Output and Prediction**: To output a prediction a^*, the model uses an output function, denoted generally as $a^* = out(\mathbf{d}, h, \mathbf{q}, g, \mathcal{A}, \mathcal{V}, \alpha)$. This function is instantiated in various ways.

   The Stanford Reader (Chen et al., 2016) computes the attention-weighted sum of the context vectors and then an inner product with each candidate answer: $\mathbf{o} = \sum_{i=1}^{|\mathbf{p}|} \alpha_i h_i$ and $a^* = \operatorname{argmax}_{a \in \mathcal{A}} W_a^\top \mathbf{o}$, where W_a is the vector for answer a. As the Stanford Reader was developed for the anonymized CNN/Daily Mail tasks, the output matrix W was very small in their experiments. However, for LAMBADA, the set of possible answers can range over the whole vocabulary and the output matrix W can be large and difficult to train. Therefore we also experiment with a modified version of the Stanford Reader with the following output function: $a^* = \operatorname{argmax}_{a \in \mathcal{A}} e(a)^\top W \mathbf{o}$, where e is the embedding function defined in the encoding step above and W is a parameter matrix.

   For the Attention Sum (Kadlec et al., 2016), Gated-Attention (Dhingra et al., 2016), and Attention-over-Attention Readers (Cui et al., 2016), the final answer is computed by: $\forall a \in \mathcal{A},\ P(a|\mathbf{d}, \mathbf{q}) = \sum_{i \in I(a,\mathbf{d})} \alpha_i$, where $I(a, \mathbf{d})$ is the set of positions where a appears in context **d**, and $a^* = \operatorname{argmax}_{a \in \mathcal{A}} P(a|\mathbf{d}, \mathbf{q})$.

[1] We overload the e function to operate on sequences and denote the embedding of **d** and **q** as matrices $e(\mathbf{d})$ and $e(\mathbf{q})$.

2.2 Training Data Construction

Each LAMBADA instance is divided into a **context** (4.6 sentences on average) and a **target sentence**, and the last word of the target sentence is the **target word** to be predicted. The LAMBADA dataset consists of development (DEV) and test (TEST) sets; Paperno et al. (2016) also provide a control dataset (CONTROL), an unfiltered sample of passages from the BookCorpus.

We construct a new training dataset from the BookCorpus. We restrict it to instances that contain the target word in the context. This decision is natural given our use of neural readers that assume the answer is contained in the passage. We also ensure that the context has at least 50 words and contains 4 or 5 sentences and we require the target sentences to have more than 10 words.

Some neural readers require a candidate target word list to choose from. We list all words in the context as candidate answers, except for punctuation symbols[2]. Our newly constructed dataset contains 1,827,123 instances in total. We further divide this dataset into two parts, a training set (TRAIN) of 1,618,782 instances and a validation set (VAL) of 208,341 instances. These datasets can be found at the authors' websites.

3 Experiments

We use the Stanford Reader (Chen et al., 2016), our modified Stanford Reader, the Attention Sum (AS) Reader (Kadlec et al., 2016), and the Gated-Attention (GA) Reader (Dhingra et al., 2016). We

[2] The list of punctuation symbols we used is available at `https://raw.githubusercontent.com/ZeweiChu/lambada-dataset/master/stopwords/shortlist-stopwords.txt`

also add the simple features from Wang et al. (2016) to the AS and GA Readers. We train each reader on TRAIN for a max of 10 epochs, stopping when accuracy on DEV decreases two epochs in a row. We take the model from the epoch with max DEV accuracy and evaluate it on TEST and CONTROL. VAL is not used.

We evaluate several other baseline systems inspired by those of Paperno et al. (2016), but we focus on versions that restrict the choice of answers to non-stopwords in the context.[3] We found this strategy to consistently improve performance even though it limits the maximum achievable accuracy.

We consider two n-gram language model baselines. We use the SRILM toolkit (Stolcke, 2002) to estimate a 4-gram model with modified Kneser-Ney smoothing on the combination of TRAIN and VAL. One uses a cache size of 100 and the other does not use a cache. We use each model to score each non-stopword from the context as the target word. We also evaluate an LSTM language model. We train the LSTM on TRAIN, where the training objective is cross entropy summed over all positions in each instance. The output vocabulary is the vocabulary of TRAIN, which is approximately 130k word types. At test time, we again limit the search to non-stopwords in the context.

We also test simple baselines that choose particular non-stopwords from the context, including a random one, the first in the context, the last in the context, and the most frequent in the context.

4 Results

Table 1 shows our results. We report accuracies on the entirety of TEST and CONTROL, as well as separately the part of CONTROL where the target word is in the context ("context"). The first part of the table shows results from Paperno et al. (2016). We then show our baselines that choose a word from the context. Choosing the most frequent yields a surprisingly high accuracy of 11.7%, which is better than all results from Paperno et al.

Our language models perform comparably, with the n-gram + cache model doing best. By forcing language models to select a word from the context, the accuracy on TEST is much higher than the analogous models from Paperno et al., though accuracy suffers on CONTROL.

We then show results with the neural readers, showing that they give much higher accuracies on TEST than all other methods. The GA Reader with simple features (Wang et al., 2016) yields the highest accuracy, reaching 49.0%.

Method	TEST	CONTROL	
	all	all	context
Baselines (Paperno et al., 2016)			
Random in context	1.6	0	N/A
Random capitalized in context	7.3	0	N/A
n-gram	0.1	19.1	N/A
n-gram + cache	0.1	19.1	N/A
LSTM	0	**21.9**	N/A
Memory network	0	8.5	N/A
Our context-restricted non-stopword baselines			
Random	5.6	0.3	2.2
First	3.8	0.1	1.1
Last	6.2	0.9	6.5
Most frequent	11.7	0.4	8.1
Our context-restricted language model baselines			
n-gram	10.7	2.2	15.6
n-gram + cache	11.8	2.2	15.6
LSTM	9.2	2.4	16.9
Our neural reader results			
Stanford Reader	21.7	7.0	49.3
Modified Stanford Reader	32.1	7.4	52.3
AS Reader	41.4	8.5	60.2
AS Reader + features	44.5	8.6	60.6
GA Reader	45.4	8.8	62.5
GA Reader + features	**49.0**	9.3	**65.6**
Human	*86.0**	-	-

Table 1: Accuracies on TEST and CONTROL datasets, computed over all instances ("all") and separately on those in which the answer is in the context ("context"). The first section is from Paperno et al. (2016). *Human accuracy is estimated from 100 randomly-sampled DEV instances.

We find that the AS and GA Readers work much better than the Stanford Reader. One cause appears to be that the Stanford Reader learns distinct embeddings for input and answer words, as discussed above. Our modified Stanford Reader, which uses only a single set of word embeddings, improves by 10.4% absolute. Since the AS and GA Readers merely score words in the context, they do not learn separate answer word embeddings and therefore do not suffer from this effect.

We suspect the remaining difference between the Stanford and the other readers is due to the difference in the output function and could be reduced by augmenting the word embeddings with generic identifiers as done in anonymization (Wang et al., 2016).

4.1 Manual Analysis

One annotator sampled 100 instances randomly from DEV, hid the final word, and attempted to guess it from the context and target sentence. The annotator was correct in 86 cases. For the subset

[3] We use the stopword list from Richardson et al. (2013).

Label	#	GA+	Human
single name cue	9	89%	100%
simple speaker tracking	19	84%	100%
basic reference	18	56%	72%
discourse inference rule	16	50%	88%
semantic trigger	20	40%	80%
coreference	21	38%	90%
external knowledge	24	21%	88%
all	100	55%	86%

Table 2: Labels derived from manual analysis of 100 LAMBADA DEV instances. An instance can be tagged with multiple labels, hence the sum of instances across labels exceeds 100.

that contained the answer in the context, the annotator was correct in 79 of 87 cases. Even though two annotators were able to give the correct answer all LAMBADA instances during dataset construction (Paperno et al., 2016), our results give an estimate of how often a third would agree.

We analyzed these 100 DEV instances, tagging each with labels indicating the minimal kinds of understanding needed to answer it correctly.[4] Each instance can have multiple labels. We briefly describe each label below:

- single name cue: the answer is clearly a name according to contextual cues and only a single name is mentioned in the context.
- simple speaker tracking: instance can be answered merely by tracking who is speaking without understanding what they are saying.
- basic reference: answer is a reference to something mentioned in the context; simple understanding/context matching suffices.
- discourse inference rule: answer can be found by applying a single discourse inference rule, such as the rule: "X left Y and went in search of Z" $\rightarrow Y \neq Z$.
- semantic trigger: amorphous semantic information is needed to choose the answer, typically related to event sequences to dialogue turns, e.g., a customer says "Where is the X?" and a supplier responds "We got plenty of X".
- coreference: instance requires non-trivial coreference resolution to solve correctly, typically the resolution of anaphoric pronouns.
- external knowledge: some particular external knowledge is needed to choose the answer.

[4]The annotations are available at the authors' websites.

Table 2 shows the breakdown of these labels across instances, as well as the accuracy on each label of the GA Reader with features.

We find that the GA Reader performs well on instances involving shallower, more surface-level cues. In 9% of cases, the answer is clearly a name based on contextual cues in the target sentence and there is only one name in the context; the reader answers all but one correctly. When only simple speaker tracking is needed (19% of cases), the reader gets 84% correct.

We find that the hardest instances are those that involve deeper understanding, like semantic links, coreference resolution, and external knowledge. While external knowledge is difficult to define, we chose this label when we were able to explicitly write down the knowledge that one would use when answering the instances, e.g., one instance requires knowing that "when something explodes, noise emanates from it". These instances make up nearly a quarter of those we analyzed, making LAMBADA a good task for work in leveraging external knowledge for language understanding.

4.2 Discussion

On CONTROL, while our readers outperform our other baselines, they are outperformed by the language modeling baselines from Paperno et al. This suggests that though we have improved the state of the art on LAMBADA by more than 40% absolute, we have not solved the general language modeling problem; there is no single model that performs well on both TEST and CONTROL.

We note that TRAIN has similar characteristics to the part of CONTROL that contains the answer in the context (the final column of Table 1). We find that the ranking of systems according to this column is similar to that in the TEST column. This suggests that our simple method of dataset creation could be used to create additional training or evaluation sets for challenging language modeling problems like LAMBADA, perhaps by combining it with baseline suppression (Onishi et al., 2016).

5 Conclusion

We constructed a new training set for LAMBADA and used it to train neural readers to improve the state of the art from 7.3% to 49%. We also provided results with several other strong baselines and included a manual evaluation in an attempt to better understand the phenomena tested by the task. Our hope is that other researchers will seek models and training regimes that simultaneously

perform well on both LAMBADA and CONTROL, with the goal of solving the general problem of language modeling.

References

Danqi Chen, Jason Bolton, and Christopher D. Manning. 2016. A thorough examination of the CNN/Daily Mail reading comprehension task. In *Proc. of ACL*.

Kyunghyun Cho, Bart van Merrienboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proc. of EMNLP*.

Yiming Cui, Zhipeng Chen, Si Wei, Shijin Wang, Ting Liu, and Guoping Hu. 2016. Attention-over-attention neural networks for reading comprehension. *arXiv preprint*.

Bhuwan Dhingra, Hanxiao Liu, William W. Cohen, and Ruslan Salakhutdinov. 2016. Gated-attention readers for text comprehension. *arXiv preprint*.

Karl Moritz Hermann, Tom Koisk, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. In *Proc. of NIPS*.

Felix Hill, Antoine Bordes, Sumit Chopra, and Jason Weston. 2016. The Goldilocks principle: Reading children's books with explicit memory representations. In *Proc. of ICLR*.

Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural Computation*, 9(8).

Rudolf Kadlec, Martin Schmid, Ondrej Bajgar, and Jan Kleindienst. 2016. Text understanding with the attention sum reader network. In *Proc. of ACL*.

Takeshi Onishi, Hai Wang, Mohit Bansal, Kevin Gimpel, and David McAllester. 2016. Who did What: A large-scale person-centered cloze dataset. In *Proc. of EMNLP*.

Denis Paperno, Germn Kruszewski, Angeliki Lazaridou, Quan Ngoc Pham, Raffaella Bernardi, Sandro Pezzelle, Marco Baroni, Gemma Boleda, and Raquel Fernndez. 2016. The LAMBADA dataset: Word prediction requiring a broad discourse context. In *Proc. of ACL*.

Matthew Richardson, Christopher JC Burges, and Erin Renshaw. 2013. MCTest: A challenge dataset for the open-domain machine comprehension of text. In *Proc. of EMNLP*.

Andreas Stolcke. 2002. SRILM-an extensible language modeling toolkit. In *Proc. of Interspeech*.

Hai Wang, Takeshi Onishi, Kevin Gimpel, and David McAllester. 2016. Emergent logical structure in vector representations of neural readers. *arXiv preprint*.

Yukun Zhu, Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proc. of ICCV*.